

# On Partial Wait-Freedom in Transactional Memory

Petr Kuznetsov<sup>1</sup>   Srivatsan Ravi<sup>2</sup>

<sup>1</sup>Télécom ParisTech

<sup>2</sup>TU Berlin

## Abstract

Transactional memory (TM) is a convenient synchronization tool that allows concurrent threads to declare sequences of instructions on shared data as speculative *transactions* with “all-or-nothing” semantics. It is known that dynamic transactional memory cannot provide *wait-free* progress in the sense that every transaction commits in a finite number of its own steps. In this paper, we explore the costs of providing wait-freedom to only a *subset* of transactions. Since most transactional workloads are believed to be read-dominated, we require that read-only transactions commit in the wait-free manner, while updating transactions are guaranteed to commit only if they run in the absence of concurrency. We show that this kind of partial wait-freedom, combined with attractive requirements like read invisibility or disjoint-access parallelism, incurs considerable complexity costs.

# 1 Introduction

The transactional memory abstraction (*TM*) allows concurrent processes to declare sequences of operations on shared data as *atomic transaction*. A transaction may *commit* in which case its updates to shared data items “take effect”, or *abort*, in which case the transaction does not affect other transactions. A *TM implementation* provides processes with algorithms for transactional operations (*read*, *write*, *tryCommit*) using low-level *base objects*.

A natural consistency criterion provided by most TM implementations that we assume in this paper is *strict serializability*: all committed transactions appear to execute sequentially in some total order respecting the timing of non-overlapping transactions. This is strictly weaker than conditions like *opacity* [10] and *virtual-world consistency* [14], which require *every* transaction (including aborted and incomplete ones) to observe a consistent state.

The spectrum of progress properties specifying the conditions under which a transaction must commit appears more interesting. Perhaps, the most attractive progress property a TM may satisfy is *wait-freedom*: every transaction commits regardless of the behavior of concurrent processes [12]. Intuitively, wait-free progress is desirable because it guarantees that a transaction takes effect in a finite number of its own steps and, thus, the progress of the corresponding process is not affected by other processes. It is easy to see, however, that dynamic TMs in which data sets of transactions are not known in advance do not allow for wait-free implementations [10]. Suppose that a transaction  $T_1$  reads data item  $X$ , then a concurrent transaction  $T_2$  reads data item  $Y$ , writes to  $X$  and commits, and finally  $T_2$  writes to  $Y$ . Since  $T_1$  has read the “old” value in  $X$  and  $T_2$  has read the “old” value in  $Y$ , there is no way to commit  $T_1$  and order the two transactions in a sequential execution. As this scenario can be repeated arbitrarily often, even the weaker guarantee of *local progress* that only requires that each transaction *eventually* commits if repeated sufficiently often, cannot be ensured by *any* strictly serializable TM implementation, regardless of the base objects it uses [4].

But can we ensure that at least *some* transactions commit wait-free? It is often argued that many realistic workloads are *read-dominated*: the proportion of read-only transactions is higher than that of updating ones, or read-only transactions have much larger data sets than updating ones [3, 11]. Therefore, it seems natural to require first that read-only transactions commit wait-free. Of course, the progress guarantees for updating transaction have to be weaker. Since this paper focuses on complexity lower bounds for read-only transactions, we consider a very weak property saying that an updating transaction is guaranteed to commit only if it runs in the absence of concurrency (we refer to this as sequential progress for updating transactions).

First, we show that if we wish to derive implementations that use *invisible reads*, i. e., assuming that read-only transactions do not apply any nontrivial primitives on the base memory, then the implementation must maintain unbounded sets of values for every data item. Since read-invisibility is believed important for (most common) read-dominated workloads, our result suggests that these implementations may not be practical due to their space complexity.

We then focus on *disjoint-access-parallel* (DAP) implementations [2, 15]. The idea of DAP is to allow transactions that do not contend on the same data item to proceed independently of each other without memory contention. A *strict DAP* TM implementation ensures that two transactions *contend* on a base object (i.e. both access the base object and at least one modifies it) only if they access a common data item [5, 9]. Interestingly, we prove that it is impossible to implement strict DAP implementations that ensure wait-free progress for read-only transactions and sequential progress for updating transactions. Thus, two transactions that access mutually disjoint data sets may prevent each other from committing.

A less restrictive definition of DAP satisfied by several popular TM implementations [8, 13, 19] is *weak DAP* [2] which is typically defined using a *conflict graph* defined for each pair of

concurrent transactions  $T_1$  and  $T_2$ . The vertices of the graph are data items accessed by  $T_1$ ,  $T_2$  and all transactions that are concurrent to them. There is an edge between two data items if the item is accessed by two concurrent transactions in the set. A weak DAP implementation ensures that transactions  $T_1$  and  $T_2$  are allowed to *concurrently contend* on a base object (i.e., to concurrently have enabled operations on it one of which is about to modify the object) only if there is a path in their conflict graph between a data item accessed by  $T_1$  and a data item accessed by  $T_2$ . For weak DAP TMs, we show that a read-only transaction (with possibly unbounded read set) must sometimes perform at least one expensive synchronization pattern [1] per read operation (the expensive-pattern complexity of a read-only transaction is linear in the size of its data set). These patterns include *read-after-write* (or *RAW*), which incurs a costly memory fence on most CPU architectures, or a *atomic-write-after-read* (or *AWAR*), typically instantiated as *atomic compare-and-swap*. The metric appears to be a more adequate complexity measure than simple step complexity, as it accounts for expensive cache-coherence operations or conditional instructions.

Overall, our results highlight considerable complexity costs incurred by requiring partial wait-freedom for read-only transactions, even when remaining updating transactions are only provided with extremely weak progress. We hope this paper provides a better understanding of the pros and cons of diversified progress guarantees for different workloads.

**Roadmap.** Section 2 describes our TM model and Section 3 defines the TM classes considered in this paper. We present in Section 4, the space complexity of implementations that use invisible reads. In Section 5, we prove the impossibility result concerning strict DAP TMs and in Section 6, the lower bound on the number of expensive synchronization patterns for weak DAP TMs. Section 7 relates our work to earlier results. In Section 8, we conclude the paper and discuss open questions.

## 2 Model

**TM interface.** A *transactional memory* (in short, *TM*) supports *transactions* for reading and writing on a finite set  $\mathcal{X}$  of data items, referred to as *t-objects*. Every transaction  $T_k$  has a unique identifier  $k$ . We assume no bound on the size of a t-object i.e. the number of possible different values a t-object can have. A transaction  $T_k$  may contain the following *t-operations*, each being a matching pair of an *invocation* and a *response*:  $read_k(X)$  returns a value in some domain  $V$  (denoted  $read_k(X) \rightarrow v$ ) or a special value  $A_k \notin V$  (*abort*);  $write_k(X, v)$ , for a value  $v \in V$ , returns *ok* or  $A_k$ ;  $tryC_k$  returns  $C_k \notin V$  (*commit*) or  $A_k$ .

**TM implementations.** We consider an asynchronous shared-memory system in which a set of  $n$  processes, communicate by applying *primitives* on shared *base objects*. We assume that processes issue transactions sequentially i.e. a process starts a new transaction only after the previous transaction has committed or aborted. A *TM implementation* provides processes with algorithms for implementing  $read_k$ ,  $write_k$  and  $tryC_k()$  of a transaction  $T_k$  by *applying primitives* from a set of shared *base objects*, each of which is assigned an *initial value*. We assume that these primitives are *deterministic*. A primitive is a generic *read-modify-write* (*RMW*) procedure applied to a base object [7, 12]. It is characterized by a pair of functions  $\langle g, h \rangle$ : given the current state of the base object,  $g$  is an *update function* that computes its state after the primitive is applied, while  $h$  is a *response function* that specifies the outcome of the primitive returned to the process. A RMW primitive is *trivial* if it never changes the value of the base object to which it is applied. Otherwise, it is *nontrivial*.

**Executions and configurations.** An *event* of a transaction  $T_k$  (sometimes we say *step* of  $T_k$ ) is an invocation or response of a t-operation performed by  $T_k$  or a RMW primitive  $\langle g, h \rangle$  applied by

$T_k$  to a base object  $b$  along with its response  $r$  (we call it a *RMW event* and write  $(b, \langle g, h \rangle, r, k)$ ).

A *configuration* (of a TM implementation) specifies the value of each base object and the state of each process. The *initial configuration* is the configuration in which all base objects have their initial values and all processes are in their initial states.

An *execution fragment* is a (finite or infinite) sequence of events. An *execution* of a TM implementation  $M$  is an execution fragment where, starting from the initial configuration, each event is issued according to  $M$  and each response of a RMW event  $(b, \langle g, h \rangle, r, k)$  matches the state of  $b$  resulting from all preceding events. An execution  $E \cdot E'$ , denoting the concatenation of  $E$  and  $E'$ , is an *extension* of  $E$  and we say that  $E'$  *extends*  $E$ .

Let  $E$  be an execution fragment. For every transaction identifier  $k$ ,  $E|k$  denotes the subsequence of  $E$  restricted to events of transaction  $T_k$ . If  $E|k$  is non-empty, we say that  $T_k$  *participates* in  $E$ , else we say  $E$  is  $T_k$ -*free*. Two executions  $E$  and  $E'$  are *indistinguishable* to a set  $\mathcal{T}$  of transactions, if for each transaction  $T_k \in \mathcal{T}$ ,  $E|k = E'|k$ .

The *read set* (resp., the *write set*) of a transaction  $T_k$  in an execution  $E$ , denoted  $Rset(T_k)$  (and resp.  $Wset(T_k)$ ), is the set of t-objects on which  $T_k$  invokes reads (and resp. writes) in  $E$ . The *data set* of  $T_k$  is  $Dset(T_k) = Rset(T_k) \cup Wset(T_k)$ . A transaction is called *read-only* if  $Wset(T_k) = \emptyset$ ; *write-only* if  $Rset(T_k) = \emptyset$  and *updating* if  $Wset(T_k) \neq \emptyset$ .

**Transaction orders.** Let  $txns(E)$  denote the set of transactions that participate in  $E$ . The *history exported by an execution*  $E$  is the subsequence of  $E$  consisting of the invocation and response events of t-operations. Two histories  $H$  and  $H'$  are *equivalent* if  $txns(H) = txns(H')$  and for every transaction  $T_k \in txns(H)$ ,  $H|k = H'|k$ . An execution  $E$  is *sequential* if every invocation of a t-operation is either the last event in the history  $H$  exported by  $E$  or is immediately followed by a matching response. We assume that executions are *well-formed* i.e. for all  $T_k$ ,  $E|k$  is sequential and has no events after  $A_k$  or  $C_k$ . A transaction  $T_k \in txns(E)$  is *complete in*  $E$  if  $E|k$  ends with a response event. The execution  $E$  is *complete* if all transactions in  $txns(E)$  are complete in  $E$ . A transaction  $T_k \in txns(E)$  is *t-complete* if  $E|k$  ends with  $A_k$  or  $C_k$ ; otherwise,  $T_k$  is *t-incomplete*.  $T_k$  is *committed* (resp., *aborted*) in  $E$  if the last event of  $T_k$  is  $C_k$  (resp.,  $A_k$ ). The execution  $E$  is *t-complete* if all transactions in  $txns(E)$  are t-complete.

For transactions  $\{T_k, T_m\} \in txns(E)$ , we say that  $T_k$  *precedes*  $T_m$  in the *real-time order* of  $E$ , denoted  $T_k \prec_E^{RT} T_m$ , if  $T_k$  is t-complete in  $E$  and the last event of  $T_k$  precedes the first event of  $T_m$  in  $E$ . If neither  $T_k \prec_E^{RT} T_m$  nor  $T_m \prec_E^{RT} T_k$ , then  $T_k$  and  $T_m$  are *concurrent* in  $E$ . An execution  $E$  is *t-sequential* if there are no concurrent transactions in  $E$ .

**Contention.** We say that a configuration  $C$  after an execution  $E$  is *quiescent* (and resp. *t-quiescent*) if every transaction  $T_k \in txns(E)$  is complete (and resp. t-complete) in  $C$ . If a transaction  $T$  is incomplete in an execution  $E$ , it has exactly one *enabled* event, which is the next event the transaction will perform according to the TM implementation. Events  $e$  and  $e'$  of an execution  $E$  *contend* on a base object  $b$  if they are both events on  $b$  in  $E$  and at least one of them is nontrivial (the event is trivial if it is the application of a trivial primitive; otherwise, nontrivial).

We say that a transaction  $T$  is *poised to apply an event*  $e$  *after*  $E$  if  $e$  is the next enabled event for  $T$  in  $E$ . We say that transactions  $T$  and  $T'$  *concurrently contend on*  $b$  *in*  $E$  if they are each poised to apply contending events on  $b$  after  $E$ .

We say that an execution fragment  $E$  is *step contention-free for t-operation*  $op_k$  if the events of  $E|op_k$  are contiguous in  $E$ . We say that an execution fragment  $E$  is *step contention-free for*  $T_k$  if the events of  $E|k$  are contiguous in  $E$ . We say that  $E$  is *step contention-free* if  $E$  is step contention-free for all transactions that participate in  $E$ .

### 3 TM classes

**TM correctness.** For a history  $H$ , a *completion* of  $H$ , denoted  $\bar{H}$ , is a history derived from  $H$  through the following procedure: (1) for every incomplete t-operation  $op_k$  of  $T_k \in \text{txns}(H)$  in  $H$ , if  $op_k = \text{read}_k \vee \text{write}_k$ , insert  $A_k$  somewhere after the invocation of  $op_k$ ; otherwise, if  $op_k = \text{try}C_k()$ , insert  $C_k$  or  $A_k$  somewhere after the last event of  $T_k$ . (2) for every complete transaction  $T_k$  that is not t-complete, insert  $\text{try}C_k \cdot A_k$  somewhere after the last event of transaction  $T_k$ .

For simplicity of presentation, we assume that each execution  $E$  begins with an “imaginary” transaction that writes initial values to all t-objects and commits before any other transaction begins in  $E$ . Let  $E$  be a t-sequential execution. For every operation  $\text{read}_k(X)$  in  $E$ , we define the *latest written value* of  $X$  as follows:

- (1) If  $T_k$  contains a  $\text{write}_k(X, v)$  preceding  $\text{read}_k(X)$ , then the latest written value of  $X$  is the value of the latest such write to  $X$ .
- (2) Otherwise, if  $E$  contains a  $\text{write}_m(X, v)$ ,  $T_m$  precedes  $T_k$ , and  $T_m$  commits in  $E$ , then the latest written value of  $X$  is the value of the latest such write to  $X$  in  $E$ . (This write is well-defined since  $E$  starts with an initial transaction writing to all t-objects.)

We say that  $\text{read}_k(X)$  is *legal* in a t-sequential execution  $E$  if it returns the *latest written value* of  $X$  in  $E$ , and  $E$  is *legal* if every  $\text{read}_k(X)$  in  $E$  that does not return  $A_k$  is legal in  $E$ .

**Definition 1.** (*Strict serializability*) A finite history  $H$  is strictly serializable if there is a legal t-complete t-sequential history  $S$ , such that (1) for any two transactions  $T_k, T_m \in \text{txns}(H)$ , if  $T_k \prec_H^{RT} T_m$ , then  $T_k$  precedes  $T_m$  in  $S$ , and (2)  $S$  is equivalent to  $\text{cseq}(\bar{H})$ , where  $\bar{H}$  is some completion of  $H$  and  $\text{cseq}(\bar{H})$  is the subsequence of  $\bar{H}$  reduced to committed transactions in  $\bar{H}$ .

We refer to  $S$  as a *serialization* of  $H$ .

**Definition 2.** (*The class  $\mathcal{RWF}$* ) A TM implementation  $M \in \mathcal{RWF}$  iff it is strictly serializable and in its every execution:

- (wait-free progress for read-only transactions) every read-only transaction commits in a finite number of its steps, and
- (sequential progress for updating transactions) every transaction running step contention-free from a t-quiet configuration, commits in a finite number of its steps.

**Invisible reads.** We say that a TM implementation  $M$  uses *invisible reads* if in every execution  $E$  of  $M$ , and every read-only transaction  $T_k \in \text{txns}(E)$ ,  $E|k$  does not contain any nontrivial events.

**Disjoint-access parallelism (DAP).** A TM implementation  $M$  is *strictly disjoint-access parallel* (*strict DAP*) if, for all executions  $E$  of  $M$ , and for all transactions  $T_i$  and  $T_j$  that participate in  $E$ ,  $T_i$  and  $T_j$  contend on a base object in  $E$  only if  $Dset(T_i) \cap Dset(T_j) \neq \emptyset$  [10].

For an execution fragment, let  $\tau_E(T_1, T_2)$  denote the set of transactions ( $T_1$  and  $T_2$  included) that are concurrent to at least one of  $T_1$  and  $T_2$  in  $E$ . Let  $G(T_1, T_2, E)$  be an undirected graph with the vertex set  $\cup_{T \in \tau_E(T_1, T_2)} Dset(T)$  and there is an edge between t-objects  $X$  and  $Y$  iff there exists  $T \in \tau_E(T_1, T_2)$  such that  $\{X, Y\} \in Dset(T)$ . We say that  $T_1$  and  $T_2$  are *disjoint-access* in  $E$  if there is no path between a t-object in  $Dset(T_1)$  and a t-object in  $Dset(T_2)$  in  $G(T_1, T_2, E)$ . A TM implementation  $M$  is *weak DAP* if, in all executions  $E$  of  $M$ , any two transactions  $T_1$  and  $T_2$  concurrently contend on the same base object after  $E$  only if  $T_1$  and  $T_2$  are not disjoint-access in  $E$  or there exists a t-object  $X \in Dset(T_1) \cap Dset(T_2)$  [2, 18].

Observe that every strict DAP TM implementation satisfies weak DAP, but not vice-versa.

We first prove the following auxiliary result, inspired by [2]:

**Lemma 1.** *Let  $M$  be any weak DAP TM implementation. Let  $\alpha \cdot \rho_1 \cdot \rho_2$  be any execution of  $M$  where  $\rho_1$  (and resp.  $\rho_2$ ) is the step contention-free execution fragment of transaction  $T_1 \notin \text{trns}(\alpha)$  (and resp.  $T_2 \notin \text{trns}(\alpha)$ ) and transactions  $T_1, T_2$  are disjoint-access in  $\alpha \cdot \rho_1 \cdot \rho_2$ . Then,  $T_1$  and  $T_2$  do not contend on any base object in  $\alpha \cdot \rho_1 \cdot \rho_2$ .*

*Proof.* Suppose, by contradiction that  $T_1$  and  $T_2$  contend on the same base object in  $\alpha \cdot \rho_1 \cdot \rho_2$ .

If in  $\rho_1$ ,  $T_1$  performs a nontrivial event on a base object on which they contend, let  $e_1$  be the last event in  $\rho_1$  in which  $T_1$  performs such an event to some base object  $b$  and  $e_2$ , the first event in  $\rho_2$  that accesses  $b$ . Otherwise,  $T_1$  only performs trivial events in  $\rho_1$  to base objects on which it contends with  $T_2$  in  $\alpha \cdot \rho_1 \cdot \rho_2$ : let  $e_2$  be the first event in  $\rho_2$  in which  $\rho_2$  performs a nontrivial event to some base object  $b$  on which they contend and  $e_1$ , the last event of  $\rho_1$  in  $T_1$  that accesses  $b$ .

Let  $\rho'_1$  (and resp.  $\rho'_2$ ) be the longest prefix of  $\rho_1$  (and resp.  $\rho_2$ ) that does not include  $e_1$  (and resp.  $e_2$ ). Since before accessing  $b$ , the execution is step contention-free for  $T_1$ ,  $\alpha \cdot \rho'_1 \cdot \rho'_2$  is an execution of  $M$ . By construction,  $T_1$  and  $T_2$  are disjoint-access in  $\alpha \cdot \rho'_1 \cdot \rho'_2$  and  $\alpha \cdot \rho_1 \cdot \rho'_2$  is indistinguishable to  $p_2$  from  $\alpha \cdot \rho'_1 \cdot \rho'_2$ . Hence,  $T_1$  and  $T_2$  are poised to apply contending events  $e_1$  and  $e_2$  on  $b$  in the configuration after  $\alpha \cdot \rho'_1 \cdot \rho'_2$ —a contradiction since  $T_1$  and  $T_2$  cannot concurrently contend on the same base object.  $\square$

## 4 On the cost of invisible reads

We prove that every TM implementation  $M \in \mathcal{RWF}$  that uses invisible reads must keep unbounded sets of values for every t-object. To do so, for every  $c \in \mathbb{N}$ , we construct an execution of  $M$  that *maintains at least  $c$  distinct values for every t-object*. We require the following technical definition:

**Definition 3.** *Let  $E$  be any execution of a TM implementation  $M$ . We say that  $E$  maintains  $c$  distinct values  $\{v_1, \dots, v_c\}$  of t-object  $X$ , if there exists an execution  $E \cdot E'$  of  $M$  such that*

- *$E'$  contains the complete executions of  $c$  t-reads of  $X$  and,*
- *for all  $i \in \{1, \dots, c\}$ , the response of the  $i^{\text{th}}$  t-read of  $X$  in  $E'$  is  $v_i$ , and if the response of the  $i^{\text{th}}$  t-read of  $X$  in  $E'$  is  $r \neq v_i$ , then  $E \cdot E'$  is not an execution of  $M$ .*

**Theorem 2.** *Let  $M$  be any TM implementation in  $\mathcal{RWF}$  that uses invisible reads, and  $\mathcal{X}$ , the set of all possible t-objects that may be accessed in any execution of  $M$ . Then, for every  $c \in \mathbb{N}$ , there exists an execution  $E$  of  $M$  such that  $E$  maintains at least  $c$  distinct values of each t-object  $X \in \mathcal{X}$ .*

*Proof.* Let  $v_{0_\ell}$  be the initial value of t-object  $X_\ell \in \mathcal{X}$ . For every  $c \in \mathbb{N}$ , we iteratively construct an execution  $E$  of  $M$  of the form depicted in Figure 1a. The construction of  $E$  proceeds in phases: there are at most  $c - 1$  phases. For all  $i \in \{0, \dots, c - 1\}$ , we denote the execution after phase  $i$  as  $E_i$  which is defined as follows:

- $E_0$  is the complete step contention-free execution fragment  $\alpha_0$  of read-only transaction  $T_0$  that performs  $\text{read}_0(X_1) \rightarrow v_{0_1}$
- for all  $i \in \{1, \dots, c - 1\}$ ,  $E_i$  is defined to be an execution of the form  $\alpha_0 \cdot \rho_1 \cdot \alpha_1 \cdots \rho_i \cdot \alpha_i$  such that for all  $j \in \{1, \dots, i\}$ ,
  - $\rho_j$  is the t-complete step contention-free execution fragment of an updating transaction  $T_{2j-1}$  that, for all  $X_\ell \in \mathcal{X}$  writes the value  $v_{j_\ell}$  and commits
  - $\alpha_j$  is the complete step contention-free execution fragment of a read-only transaction  $T_{2j}$  that performs  $\text{read}_{2j}(X_1) \rightarrow v_{j_1}$

Since read-only transactions are invisible, for all  $i \in \{0, \dots, c - 1\}$ , the execution fragment  $\alpha_i$  does not contain any nontrivial events. Consequently, for all  $i < j \leq c - 1$ , the configuration




Figure 1: Executions in the proof of Theorem 2; execution in 1a must maintain  $c$  distinct values of every t-object

after  $E_i$  is indistinguishable to transaction  $T_{2j-1}$  from a t-quiescent configuration and it must be committed in  $\rho_j$  (by sequential progress for updating transactions). Observe that, for all  $1 \leq j < i$ ,  $T_{2j-1} \prec_E^{RT} T_{2i-1}$ . Strict serializability of  $M$  now stipulates that, for all  $i \in \{1, \dots, c-1\}$ , the t-read of  $X_1$  performed by transaction  $T_{2i}$  in the execution fragment  $\alpha_i$  must return the value  $v_{i_1}$  of  $X_1$  as written by transaction  $T_{2i-1}$  in the execution fragment  $\rho_i$  (in any serialization,  $T_{2i-1}$  is the latest committed transaction writing to  $X_1$  that precedes  $T_{2i}$ ). Thus,  $M$  indeed has an execution  $E$  of the form depicted in Figure 1a.

Consider the execution fragment  $E'$  that extends  $E$  in which, for all  $i \in \{0, \dots, c-1\}$ , read-only transaction  $T_{2i}$  is extended with the complete execution of the t-reads of every t-object  $X_\ell \in \mathcal{X} \setminus \{X_1\}$  (depicted in Figure 1b).

We claim that, for all  $i \in \{0, \dots, c-1\}$ , and for all  $X_\ell \in \mathcal{X} \setminus \{X_1\}$ ,  $read_{2i}(X_\ell)$  performed by transaction  $T_{2i}$  must return the value  $v_{i_\ell}$  of  $X_\ell$  written by transaction  $T_{2i-1}$  in the execution fragment  $\rho_i$ . Indeed, by wait-free progress,  $read_i(X_\ell)$  must return a non-abort response in such an extension of  $E$ . Suppose by contradiction that  $read_i(X_\ell)$  returns a response that is not  $v_{i_\ell}$ . There are two cases:

- $read_{2i}(X_\ell)$  returns the value  $v_{j_\ell}$  written by transaction  $T_{2j-1}$ ;  $j < i$ . However, since for all  $j < i$ ,  $T_{2j} \prec_E^{RT} T_{2i}$ , the execution is not strictly serializable—contradiction.
- $read_{2i}(X_\ell)$  returns the value  $v_{j_\ell}$  written by transaction  $T_{2j}$ ;  $j > i$ . Since  $read_i(X_1)$  returns the value  $v_{i_1}$  and  $T_{2i} \prec_E^{RT} T_{2j}$ , there exists no such serialization—contradiction.

Thus,  $E$  maintains at least  $c$  distinct values of every t-object  $X \in \mathcal{X}$ .  $\square$




Figure 2: Executions in the proof of Theorem 3; execution in 2c is not strictly serializable

## 5 Impossibility of strict disjoint-access parallelism

In this section, we prove that it is impossible to derive TM implementations in  $\mathcal{RW}\mathcal{F}$  which ensure that any two transactions accessing pairwise disjoint data sets can execute without contending on the same base object.

**Theorem 3.** *There exists no strict DAP TM implementation in  $\mathcal{RW}\mathcal{F}$ .*

*Proof.* Suppose by contradiction that there exists a strict DAP TM implementation  $M \in \mathcal{RW}\mathcal{F}$ .

Let  $v$  be the initial value of t-objects  $X_1$ ,  $X_2$  and  $X_3$ . Let  $\pi$  be the t-complete step contention-free execution of transaction  $T_1$  that writes the value  $nv \neq v$  to t-objects  $X_1$  and  $X_3$ . By sequential progress for updating transactions,  $T_1$  must be committed in  $\pi$ .

Note that any read-only transaction that runs step contention-free after some prefix of  $\pi$  must return a non-abort value. Since any such transaction must return  $v$  after the empty prefix of  $\pi$  and  $nv$  when it starts from  $\pi$ , there exists  $\pi'$ , the longest prefix of  $\pi$  that cannot be extended with the t-complete step contention-free execution of any transaction that performs a t-read of  $X_1$  and returns  $nv$  nor with the t-complete step contention-free execution of any transaction that performs a t-read of  $X_3$  and returns  $nv$ .

Consider the execution fragment  $\pi' \cdot \alpha_1$ , where  $\alpha_1$  is the complete step contention-free execution of transaction  $T_0$  that performs  $read_0(X_1) \rightarrow v$ . Indeed, by definition of  $\pi'$  and wait-free progress (assumed for read-only transactions),  $M$  has an execution of the form  $\pi' \cdot \alpha_1$ .

Let  $e$  be the enabled event of transaction  $T_1$  in the configuration after  $\pi'$ . Without loss of generality, assume that  $\pi' \cdot e$  can be extended with the t-complete step contention-free execution of a transaction that reads  $X_3$  and returns  $nv$ .

We now prove that  $M$  has an execution of the form  $\pi' \cdot \alpha_1 \cdot e \cdot \beta \cdot \gamma$ , where



- $\beta$  is the t-complete step contention-free execution fragment of transaction  $T_3$  that performs  $read_3(X_3) \rightarrow nv$  and commits
- $\gamma$  is the t-complete step contention-free execution fragment of transaction  $T_2$  that writes  $nv$  to  $X_2$  and commits.

Observe that, by definition of  $\pi'$ ,  $M$  has an execution of the form  $\pi' \cdot e \cdot \beta$ . By construction, transaction  $T_1$  applies a nontrivial primitive to a base object, say  $b$  in the event  $e$  that is accessed by transaction  $T_3$  in the execution fragment  $\beta$ . Since transactions  $T_0$  and  $T_3$  access mutually disjoint data sets in  $\pi' \cdot \alpha_1 \cdot e \cdot \beta$ ,  $T_3$  does not access any base object in  $\beta$  to which transaction  $T_0$  applies a nontrivial primitive in the execution fragment  $\alpha_1$  (assumption of strict DAP). Thus,  $\alpha_1$  does not contain a nontrivial primitive to  $b$  and  $\pi' \cdot \alpha_1 \cdot e \cdot \beta$  is indistinguishable to  $T_3$  from the execution  $\pi' \cdot e \cdot \beta$ . This proves that  $M$  has an execution of the form  $\pi' \cdot \alpha_1 \cdot e \cdot \beta$  (depicted in Figure 2a).

Since transaction  $T_2$  writes to t-object  $Dset(T_2) = X_2 \notin \{Dset(T_1) \cup Dset(T_0) \cup Dset(T_3)\}$ , by strict DAP, the configuration after  $\pi' \cdot \alpha_1 \cdot e \cdot \beta$  is indistinguishable to  $T_2$  from a t-quiet configuration. Indeed, transaction  $T_2$  does not contend with any of the transactions  $T_1$ ,  $T_0$  and  $T_3$  on any base object in  $\pi' \cdot \alpha_1 \cdot e \cdot \beta \cdot \gamma$ . Sequential progress of  $M$  requires that  $T_2$  must be committed in  $\pi' \cdot \alpha_1 \cdot e \cdot \beta \cdot \gamma$ . Thus,  $M$  has an execution of the form  $\pi' \cdot \alpha_1 \cdot e \cdot \beta \cdot \gamma$ .

By the above arguments, the execution  $\pi' \cdot \alpha_1 \cdot e \cdot \beta \cdot \gamma$  is indistinguishable to each of the transactions  $T_1$ ,  $T_0$ ,  $T_2$  and  $T_3$  from  $\gamma \cdot \pi' \cdot \alpha_1 \cdot e \cdot \beta$  in which transaction  $T_2$  precedes  $T_1$  in real-time ordering. Thus,  $\gamma \cdot \pi' \cdot \alpha_1 \cdot e \cdot \beta$  is also an execution of  $M$ .

Consider the extension of the execution  $\gamma \cdot \pi' \cdot \alpha_1 \cdot e \cdot \beta$  in which transaction  $T_0$  performs  $read_0(X_2)$  and commits (depicted in Figure 2b). Strict serializability of  $M$  stipulates that  $read_0(X_2)$  must return  $nv$  since  $T_2$  (which writes  $nv$  to  $X_2$  in  $\gamma$ ) precedes  $T_0$  in this execution.

Similarly, we now extend the execution  $\pi' \cdot \alpha_1 \cdot e \cdot \beta \cdot \gamma$  with the complete step contention-free execution fragment of the t-read of  $X_2$  by transaction  $T_0$ . Since  $T_0$  is a read-only transaction, it must be committed in this extension. However, as proved above, this execution is indistinguishable to  $T_0$  from the execution depicted in Figure 2b in which  $read_0(X_2)$  must return  $nv$ . Thus,  $M$  has an execution of the form  $\pi' \cdot \alpha_1 \cdot e \cdot \beta \cdot \gamma \cdot \alpha_2$ , where  $T_0$  performs  $read_0(X_2) \rightarrow nv$  in  $\alpha_2$  and commits.

However, the execution  $\pi' \cdot \alpha_1 \cdot e \cdot \beta \cdot \gamma \cdot \alpha_2$  (depicted in Figure 2c) is not strictly serializable. Indeed, transaction  $T_1$  must be committed in any serialization and must precede transaction  $T_3$  since  $read_3(X_3)$  returns the value of  $X_3$  written by  $T_m$ . However, transaction  $T_0$  must precede  $T_1$  since  $read_0(X_1)$  returns the initial the value of  $X_1$ . Also, transaction  $T_2$  must precede  $T_0$  since  $read_0(X_2)$  returns the value of  $X_2$  written by  $T_2$ . But transaction  $T_3$  must precede  $T_2$  to respect real-time ordering of transactions. Thus,  $T_1$  must precede  $T_0$  in any serialization. But there exists no such serialization: a contradiction to the assumption that  $M$  is strictly serializable.  $\square$

## 6 A linear lower bound on expensive synchronization

Attiya *et al.* identified two common expensive synchronization patterns that frequently arise in the design of concurrent algorithms: *read-after-write (RAW)* or *atomic write-after-read (AWAR)* [1,17]. In this section, we prove a linear lower bound (in the size of the transaction's read set) on the number of RAWs or AWARs for weak DAP TM implementations in  $\mathcal{RWF}$ . To so so, we construct an execution in which each t-read operation of an arbitrarily long read-only transaction contains a RAW or an AWAR.

**Definition 4.** (*RAW/AWAR metric*) Let  $\pi$  be a fragment of an execution of a TM implementation  $M$  and let  $\pi^i$  denote the  $i$ -th event in  $\pi$  ( $i = 0, \dots, |\pi| - 1$ ).

We say that a transaction  $T$  performs a RAW (read-after-write) in  $\pi$  if  $\exists i, j; 0 \leq i < j < |\pi|$  such that (1)  $\pi^i$  is a write to a base object  $b$  by  $T$ , (2)  $\pi^j$  is a read of a base object  $b' \neq b$  by  $T$  and




Figure 3: Executions in the proof of Theorem 4; execution in 3c is not strictly serializable

(3) there is no  $\pi^k$  such that  $i < k < j$  and  $\pi^k$  is a write to  $b'$  by  $T$ .

We say a transaction  $T$  performs an AWAR (atomic-write-after-read) in  $\pi$  if  $\exists i, 0 \leq i < |\pi|$  such that the event  $\pi^i$  is the application of a nontrivial rmw primitive that reads a base object  $b$  followed by a write to  $b$ .

**Theorem 4.** Every weakly DAP TM implementation  $M \in \mathcal{RWF}$  has, for all  $m \in \mathbb{N}$ , an execution in which some read-only transaction  $T_0$  with  $m = |Rset(T_0)|$  performs  $\Omega(m)$  RAWs/AWARs.

*Proof.* Let  $v$  be the initial value of each of the t-objects  $X_1, \dots, X_m$ . Consider the t-complete step contention-free execution of transaction  $T_0$  that performs  $m$  t-reads  $read_0(X_1), read_0(X_1), \dots, read_0(X_m)$  and commits. We prove that each of the first  $m - 1$  t-reads must perform a RAW or an AWAR.

For each  $j \in \{1, \dots, m - 1\}$ , we define an execution of the form  $\alpha_1 \cdot \alpha_2 \cdots \alpha_j$ , where for all  $i \in \{1, \dots, j\}$ ,  $\alpha_i$  is the complete step contention-free execution fragment of  $read_0(X_j) \rightarrow v$ . Suppose by contradiction that  $\alpha_j$  does not contain a RAW or an AWAR.

The following claim shows that we can schedule a committed transaction  $T_j$  that writes a new value to  $X_j$  concurrent to  $read_0(X_j)$  such that the execution is indistinguishable to both  $T_0$  and  $T_j$  from a step contention-free execution (depicted in Figure 3a).

**Claim 5.** For all  $j \in \{1, \dots, m - 1\}$ ,  $M$  has an execution of the form  $\alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^1 \cdot \delta_j \cdot \alpha_j^2$  where,

- $\delta_j$  is the t-complete step contention-free execution fragment of transaction  $T_j$  that writes  $nv \neq v$  and commits
- $\alpha_j^1 \cdot \alpha_j^2 = \alpha_j$  is the complete execution fragment of the  $j^{th}$  t-read  $read_0(X_j) \rightarrow v$  such that
  - $\alpha_j^1$  does not contain any nontrivial events

–  $\alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^1 \cdot \delta_j \cdot \alpha_j^2$  is indistinguishable to  $T_0$  from the step contention-free execution fragment  $\alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^1 \cdot \alpha_j^2$

Moreover,  $T_j$  does not access any base object to which  $T_0$  applies a nontrivial event in  $\alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^1 \cdot \delta_j$ .

*Proof.* By wait-free progress (for read-only transactions) and strict serializability,  $M$  has an execution of the form  $\alpha_1 \cdots \alpha_{j-1}$  in which each of the t-reads performed by  $T_0$  must return the initial value of the t-objects.

Since  $T_j$  is an updating transaction, by sequential progress, there exists an execution of  $M$  of the form  $\delta_j \cdot \alpha_1 \cdots \alpha_{j-1}$ . Since  $T_0$  and  $T_j$  are disjoint-access in the  $\delta_j \cdot \alpha_1 \cdots \alpha_{j-1}$ , by Lemma 1,  $T_0$  and  $T_j$  do not contend on any base object in  $\delta_j \cdot \alpha_1 \cdots \alpha_{j-1}$ . Thus,  $\alpha_1 \cdots \alpha_{j-1} \cdot \delta_j$  is indistinguishable to  $T_j$  from the execution  $\delta_j$  and  $\alpha_1 \cdots \alpha_{j-1} \cdot \delta_j$  is also an execution of  $M$ .

Let  $e$  be the first event that contains a write to a base object in  $\alpha_j$ . If there exists no such write event to a base object in  $\alpha_j$ , then  $\alpha_j^1 = \alpha_j$  and  $\alpha_j^2$  is empty. Otherwise, we represent the execution fragment  $\alpha_j$  as  $\alpha_j^1 \cdot e \cdot \alpha_j^f$ .

Since  $\alpha_j^s$  does not contain any nontrivial events that write to a base object,  $\alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^s \cdot \delta_j$  is indistinguishable to transaction  $T_j$  from the execution  $\alpha_1 \cdots \alpha_{j-1} \cdot \delta_j$ . Thus,  $\alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^s \cdot \delta_j$  is an execution of  $M$ . Since  $e$  is not an atomic-write-after-read,  $\alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^s \cdot \delta_j \cdot e$  is an execution of  $M$ . Since  $\alpha_j$  does not contain a RAW, any read performed in  $\alpha_j^f$  may only be performed to base objects previously written in  $e \cdot \alpha_j^f$ . Thus,  $\alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^s \cdot \delta_j \cdot e \cdot \alpha_j^f$  is indistinguishable to transaction  $T_0$  from the step contention-free execution  $\alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^s \cdot e \cdot \alpha_j^f$  in which  $read_0(X_j) \rightarrow v$ .

Choosing  $\alpha_j^2 = e \cdot \alpha_j^f$ , it follows that  $M$  has an execution of the form  $\alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^1 \cdot \delta_j \cdot \alpha_j^2$  that is indistinguishable to  $T_j$  and  $T_0$  from a step contention-free execution. The proof follows.  $\square$

We now prove that, for all  $j \in \{1, \dots, m-1\}$ ,  $M$  has an execution of the form  $\delta_m \cdot \alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^1 \cdot \delta_j \cdot \alpha_j^2$  such that

- $\delta_m$  is the t-complete step contention-free execution of transaction  $T_\ell$  that writes  $nv \neq v$  to  $X_m$  and commits
- $T_\ell$  and  $T_0$  do not contend on any base object in  $\delta_m \cdot \alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^1 \cdot \delta_j \cdot \alpha_j^2$
- $T_\ell$  and  $T_j$  do not contend on any base object in  $\delta_m \cdot \alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^1 \cdot \delta_j \cdot \alpha_j^2$ .

By sequential progress for updating transactions,  $T_\ell$  which writes the value  $nv$  to  $X_m$  must be committed in  $\delta_m$  since it is running in the absence of step-contention from the initial configuration. Observe that  $T_\ell$  and  $T_0$  are disjoint-access in  $\delta_m \cdot \alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^1 \cdot \delta_j \cdot \alpha_j^2$ . By definition of  $\alpha_j^1$  and  $\alpha_j^2$ ,  $\delta_m \cdot \alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^1 \cdot \delta_j \cdot \alpha_j^2$  is indistinguishable to  $T_0$  from  $\delta_m \cdot \alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^1 \cdot \alpha_j^2$ . By Lemma 1,  $T_\ell$  and  $T_0$  do not contend on any base object in  $\delta_m \cdot \alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^1 \cdot \alpha_j^2$ .

By Claim 5,  $\delta_m \cdot \alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^1 \cdot \delta_j$  is indistinguishable to  $T_j$  from  $\delta_m \cdot \delta_j$ . But transactions  $T_\ell$  and  $T_j$  are disjoint-access in  $\delta_m \cdot \delta_j$ , and by Lemma 1,  $T_j$  and  $T_\ell$  do not contend on any base object in  $\delta_m \cdot \delta_j$ .

Since strict serializability of  $M$  stipulates that each of the  $j$  t-reads performed by  $T_0$  return the initial values of the respective t-objects,  $M$  has an execution of the form  $\delta_m \cdot \alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^1 \cdot \delta_j \cdot \alpha_j^2$ .

Consider the extension of  $\delta_m \cdot \alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^1 \cdot \delta_j \cdot \alpha_j^2$  in which  $T_0$  performs  $(m-j)$  t-reads of  $X_{j+1}, \dots, X_m$  step contention-free and commits (depicted in Figure 3b). By wait-free progress of  $M$  and since  $T_0$  is a read-only transaction, there exists such an execution. Notice that the  $m^{th}$  t-read,  $read_0(X_m)$  must return the value  $nv$  by strict serializability since  $T_\ell$  precedes  $T_0$  in real-time order in this execution.

Recall that neither pairs of transactions  $T_\ell$  and  $T_j$  nor  $T_\ell$  and  $T_0$  contend on any base object in the execution  $\delta_m \cdot \alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^1 \cdot \delta_j \cdot \alpha_j^2$ . It follows that for all  $j \in \{1, \dots, m-1\}$ ,  $M$  has an execution of the form  $\alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^1 \cdot \delta_j \cdot \alpha_j^2 \cdot \delta_m$  in which  $T_j$  precedes  $T_\ell$  in real-time order.

Let  $\alpha'$  be the execution fragment that extends  $\alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^1 \cdot \delta_j \cdot \alpha_j^2 \cdot \delta_m$  in which  $T_0$  performs  $(m-j)$  t-reads of  $X_{j+1}, \dots, X_m$  step contention-free and commits (depicted in Figure 3c). Since  $\alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^1 \cdot \delta_j \cdot \alpha_j^2 \cdot \delta_m$  is indistinguishable to  $T_0$  from the execution  $\delta_m \cdot \alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^1 \cdot \delta_j \cdot \alpha_j^2$ ,  $read_0(X_m)$  must return the response value  $nv$  in  $\alpha'$ .

The execution  $\alpha_1 \cdots \alpha_{j-1} \cdot \alpha_j^1 \cdot \delta_j \cdot \alpha_j^2 \cdot \delta_m \cdot \alpha'$  is not strictly serializable. In any serialization,  $T_j$  must precede  $T_\ell$  to respect the real-time ordering of transactions, while  $T_\ell$  must precede  $T_0$  since  $read_j(X_m)$  returns the value of  $X_m$  updated by  $T_\ell$ . Also, transaction  $T_0$  must precede  $T_j$  since  $read_0(X_j)$  returns the initial value of  $X_j$ . But there exists no such serialization: a contradiction to the assumption that  $M$  is strict serializable.

Thus, for all  $j \in \{1, \dots, m-1\}$ , transaction  $T_0$  must perform a RAW or an AWAR during the execution of  $read_0(X_j)$ , completing the proof.  $\square$

Since Theorem 4 implies that read-only transactions must perform nontrivial events, we have the following corollary that was proved directly in [2].

**Corollary 6** ([2]). *There does not exist any weak DAP TM implementation  $M \in \mathcal{RWF}$  that uses invisible reads.*

## 7 Related work

Strict DAP was introduced by Guerraoui and Kapalka [10] who proved that it is impossible to implement *obstruction-free* TMs (transactions running in the absence of step contention must commit) that satisfy strict DAP.  $\mathcal{RWF}$  is incomparable to the class of obstruction-free TMs, as is the proof technique used to establish the impossibility.

Attia *et al.* [2] introduced the notion of weak DAP and showed that it is impossible to implement weak DAP strictly serializable TMs in  $\mathcal{RWF}$  if read-only transactions may only apply trivial primitives to base objects. Attia *et al.* [2] also considered a stronger “disjoint-access” property, called simply DAP, referring to the original definition proposed Israeli and Rappoport [15]. In DAP, two transactions are allowed to *concurrently access* (even for reading) the same base object only if they are disjoint-access. For an  $n$ -process DAP TM implementation, [2] showed that a read-only transaction must perform at least  $n-3$  writes. Our lower bound is strictly stronger than the one in [2], as it assumes only weak DAP, considers a more precise RAW/AWAR metric, and does not depend on the number of processes in the system. (Technically, the last point follows from the fact that the execution constructed in the proof of Theorem 4 uses only 3 concurrent processes.) Thus, the theorem strengthens and subsumes the two lower bounds of [2] within a single proof.

Perelman *et al.*, considered the class of *mv-permissive* TMs: only updating transactions that read-write conflict on data items with another updating transaction may be aborted [18]. In general,  $\mathcal{RWF}$  is incomparable with the class of *mv-permissive* TMs. On the one hand, as we observed, mv-permissiveness does not guarantee that read-only transactions are not guaranteed to commit in a wait-free manner. On the other hand,  $\mathcal{RWF}$  allows an updating transaction to abort in the presence of a concurrent read-only transaction, which is disallowed by mv-permissive TMs. Assuming that every t-operation returns in a finite number of its own steps (possibly with an *abort*, we call this wait-free TM-liveness), [18] showed that implementing a weak DAP mv-permissive TM is impossible. While  $\mathcal{RWF}$  is not subject to this impossibility, it is impossible to implement strict

DAP implementations in  $\mathcal{RWF}$ . In principle, the class of TMs considered in [18] is a proper subset of our  $\mathcal{RWF}$ .

Also, [18] proved that mv-permissive TMs cannot be *online space optimal* i.e. no mv-permissive TM can be optimal in terms of number of versions kept. Our result on the space complexity of implementations in  $\mathcal{RWF}$  that use invisible reads is strictly stronger since it applies to a larger class of TMs and proves that such implementations must maintain unbounded number of versions of every data item. Moreover, it is easy to see that our proof is also applicable for the class of TMs considered in [18].

Attia *et al.* introduced the RAW/AWAR metric and proved that it is impossible to derive RAW/AWAR-free implementations of a wide class of data types. The metric has been used in [16] to measure the complexity of read-only transactions in a strictly stronger class of *permissive* TMs that provide wait-free TM-liveness, in which a transaction may be aborted only if committing it would violate opacity. Detailed coverage on memory fences and the RAW/AWAR metric can be found in [17].

## 8 Concluding remarks

In this paper, we studied the issue of providing different progress guarantees to different classes of transactions, assuming that read-only transactions are wait-free, but updating transactions are guaranteed to commit only when they run sequentially. First, we prove that if read-only transactions are required to be invisible, then any strictly serializable TM implementation with these progress guarantees must maintain an unbounded number of values for every data item. Second, we prove that strictly serializable TMs with these progress guarantees cannot be disjoint-access parallel in a strict sense. Then, assuming a weaker form of disjoint-access-parallelism, we show that there exists an arbitrarily long read-only transaction that performs an expensive synchronization pattern within each of its read operations. Our lower bounds also hold for stronger TM-correctness conditions like *opacity* [10], *virtual-world consistency* [14], *TMS1* and *TMS2* [6].

Some questions remain open. Is the lower bound of Theorem 4 tight? We conjecture that it is not. Can we establish a fundamental tradeoff between the complexity of read-only and updating transactions incurred by implementations in  $\mathcal{RWF}$ ? More generally, assuming transactional operations provide wait-free termination, what kind of transactions can be provided with unconditional progress? Addressing these questions is ongoing and future work.

## References

- [1] H. Attiya, R. Guerraoui, D. Hendler, P. Kuznetsov, M. Michael, and M. Vechev. Laws of order: Expensive synchronization in concurrent algorithms cannot be eliminated. In *POPL*, pages 487–498, 2011.
- [2] H. Attiya, E. Hillel, and A. Milani. Inherent limitations on disjoint-access parallel implementations of transactional memory. *Theory of Computing Systems*, 49(4):698–719, 2011.
- [3] H. Attiya and A. Milani. Transactional scheduling for read-dominated workloads. In *Proceedings of the 13th International Conference on Principles of Distributed Systems*, OPODIS ’09, pages 3–17, Berlin, Heidelberg, 2009. Springer-Verlag.

- [4] V. Bushkov, R. Guerraoui, and M. Kapalka. On the liveness of transactional memory. In *Proceedings of the 2012 ACM Symposium on Principles of Distributed Computing, PODC '12*, pages 9–18, New York, NY, USA, 2012. ACM.
- [5] D. Dice and N. Shavit. What really makes transactions fast? In *Transact*, 2006.
- [6] S. Doherty, L. Groves, V. Luchangco, and M. Moir. Towards formally specifying and verifying transactional memory. *Formal Aspects of Computing*, 25(5):769–799, 2013.
- [7] F. Ellen, D. Hendler, and N. Shavit. On the inherent sequentiality of concurrent objects. *SIAM J. Comput.*, 41(3):519–536, 2012.
- [8] K. Fraser. Practical lock-freedom. Technical report, Cambridge University Computer Laboratory, 2003.
- [9] R. Guerraoui and M. Kapalka. On obstruction-free transactions. In *Proceedings of the twentieth annual symposium on Parallelism in algorithms and architectures, SPAA '08*, pages 304–313, New York, NY, USA, 2008. ACM.
- [10] R. Guerraoui and M. Kapalka. *Principles of Transactional Memory, Synthesis Lectures on Distributed Computing Theory*. Morgan and Claypool, 2010.
- [11] R. Guerraoui, M. Kapalka, and J. Vitek. Stmbench7: A benchmark for software transactional memory. *SIGOPS Oper. Syst. Rev.*, 41(3):315–324, Mar. 2007.
- [12] M. Herlihy. Wait-free synchronization. *ACM Trans. Prog. Lang. Syst.*, 13(1):123–149, 1991.
- [13] M. Herlihy, V. Luchangco, M. Moir, and W. N. Scherer, III. Software transactional memory for dynamic-sized data structures. In *PODC*, pages 92–101, 2003.
- [14] D. Imbs, J. R. G. de Mendivil, and M. Raynal. Brief announcement: virtual world consistency: a new condition for stm systems. In *PODC*, pages 280–281, 2009.
- [15] A. Israeli and L. Rappoport. Disjoint-access-parallel implementations of strong shared memory primitives. In *PODC*, pages 151–160, 1994.
- [16] P. Kuznetsov and S. Ravi. On the cost of concurrency in transactional memory. In *OPODIS*, pages 112–127, 2011. full version: <http://arxiv.org/abs/1103.1302>.
- [17] P. E. McKenney. Memory barriers: a hardware view for software hackers. Linux Technology Center, IBM Beaverton, June 2010.
- [18] D. Perelman, R. Fan, and I. Keidar. On maintaining multiple versions in STM. In *PODC*, pages 16–25, 2010.
- [19] F. Tabbà, M. Moir, J. R. Goodman, A. W. Hay, and C. Wang. Nztm: Nonblocking zero-indirection transactional memory. In *Proceedings of the Twenty-first Annual Symposium on Parallelism in Algorithms and Architectures, SPAA '09*, pages 204–213, New York, NY, USA, 2009. ACM.